

# Compressed optimization of device architectures

Adam Frees,<sup>1,2,3</sup> John King Gamble,<sup>1,3,\*</sup> Daniel R. Ward,<sup>2,3</sup> Robin Blume-Kohout,<sup>3</sup> M. A. Eriksson,<sup>2</sup> Mark Friesen,<sup>2</sup> and S. N. Coppersmith<sup>2</sup>

<sup>1</sup>These authors contributed equally to this work.

<sup>2</sup>Department of Physics, University of Wisconsin-Madison, Madison, WI 53706

<sup>3</sup>Center for Computing Research, Sandia National Laboratories, Albuquerque, NM 87123

Recent advances in nanotechnology have enabled researchers to control individual quantum mechanical objects with unprecedented accuracy, opening the door for both quantum<sup>1–8</sup> and extreme-scale conventional computing applications<sup>9</sup>. As these devices become larger and more complex, the ability to design them such that they can be simply controlled becomes a daunting and computationally infeasible task<sup>10</sup>. Here, motivated by ideas from compressed sensing<sup>11,12</sup>, we introduce a protocol for the Compressed Optimization of Device Architectures (CODA). It leads naturally to a metric for benchmarking device performance and optimizing device designs, and provides a scheme for automating the control of gate operations and reducing their complexity. Because CODA is computationally efficient, it is readily extensible to large systems. We demonstrate the CODA benchmarking and optimization protocols through simulations of up to eight quantum dots in devices that are currently being developed experimentally for quantum computation.

Nanoscale devices are challenging to control in part because their size makes them susceptible to even the smallest material defects. Quantum devices are especially challenging because their energy spectra and tunnel couplings both require fine tuning<sup>13</sup>. Here we focus on quantum bits (qubits) formed of electron spins in electrostatically-gated quantum dots<sup>14</sup> and donors<sup>15</sup>. Such devices are extensible to large arrays by leveraging the mature semiconductor processing industry. Working in a variety of systems, researchers have already demonstrated complete control and excellent decoherence properties of devices with up to four quantum dots<sup>1–5,7,8,16–25</sup>, as well as impurities coupled to electronic reservoirs<sup>6,9,26–31</sup>.

Processing quantum information in semiconductors requires fine control of the energies and tunnel rates of the individual electrons. Voltages are simultaneously tuned on many electrodes to precisely shape the electrostatic potential landscape within a device. Unfortunately, gate geometries that work well for 1-2 qubits are not necessarily well suited for larger arrays. For example, electrodes designed to control a given dot will also affect its neigh-

bors via capacitive crosstalk. While voltage compensation methods are used to manage crosstalk in smaller devices, hand tuning becomes impractical for larger arrays. Recent methods have been proposed to help automate this process, including optimized randomized benchmarking for immediate tune-up (ORBIT)<sup>32</sup>, real-time Hamiltonian estimation<sup>33</sup>, and computer-automated device tuning<sup>34</sup>. While promising, these schemes are either software-based or are only directly applicable to specific device designs. In either case, these schemes are constrained by the capabilities of the quantum device hardware.

In this paper, we aim to make the control problem more tractable through hardware optimization. We show how to modify systematically dot properties, such as occupations, energies, and tunnel rates, while changing as few voltages as possible – a strategy we refer to as *control sparsity*. The scheme relies on results and methods used for compressed sensing<sup>11,12</sup> in the signal processing literature, in which the decoding of signals can be made much more efficient by exploiting their sparseness. To demonstrate these concepts, we implement the CODA protocol using realistic simulations of several quantum dot devices. We show that CODA yields solutions that are simultaneously sparse and spatially localized near the relevant dot – an extremely desirable property for extensibility. Moreover, formulating control as an optimization problem allows us to directly compare the effectiveness of different device architectures, enabling optimization of the device designs themselves.

For qubit applications, typical properties we wish to control are the quantum dot occupations and the tunnel rates between the dots, although other properties may also be of interest. Such properties are referred to as *operational targets*. Generally, an operational target could be composed of several physical attributes; we therefore represent it as a vector  $\mathbf{t}$  in a vector space of outputs  $\mathcal{T}$ . The outputs are determined by the voltages applied to the electrodes and by the device structure. A given set of voltages is referred to as a *control setting*, represented as a vector  $\mathbf{c}$  in a vector space of controls  $\mathcal{C}$ . A physical system is then representable by a function that maps the controls to the targets:  $\hat{S} : \mathcal{C} \rightarrow \mathcal{T}$ , shown in Fig. 1. Although this mapping is generally nonlinear, for the moment we limit ourselves to a neighborhood in which the mapping is linear, centered around a working point.

In an experiment, it is often the case that one particular operational target needs to be pulsed rapidly, while all

\* Correspondence should be addressed to: jkgambl@sandia.gov




FIG. 1. Illustration of the CODA procedure for characterizing and optimizing nanoscale device designs. As an example, we consider a Si/SiGe electrostatically defined quantum dot qubit device<sup>8,18</sup>. Metal electrodes on the device surface are used to accumulate and control a two-dimensional electron gas (2DEG) in the Si quantum well, forming a double quantum dot. **a**, Device simulations determine the response of the operational targets to the physical controls. The gate voltages are adjusted to tune the device to an initial working point. Then, the configuration is probed by varying the control gate voltages and noting the response of target variables of interest (e.g. quantum dot occupations and tunnel couplings). **b**, A cross-covariance matrix describing the linear response of the system about the working point is determined using the results of (a). **c**, Since there are more controls than targets, the inverse problem yields a subset of valid control vectors for every operational target. By imposing  $L_1$  regularization, we can identify control vectors that are both local and sparse. Global optimization is achieved by using predictions from a single step in the CODA cycle to identify a new working point, and iterating the procedure.

the others are held fixed. Since many different  $\mathbf{c}$  can yield a desired  $\mathbf{t}$ , we need both a metric that defines the best  $\mathbf{c}$  and an efficient way to find this control. It is intuitive that solutions are not all equivalently useful: some have extreme voltages that are not experimentally practical, some involve precisely manipulating all the control voltages of the system, while others involve only modest voltages on a few electrodes. The latter scenario describes a desirable property often referred to as “orthogonal” control; *i.e.*, that relatively few electrodes should be involved when trying to change an operational target. Hence, our goal is to identify solutions in which the resulting change in applied voltages is as simple as possible to implement.

This problem is exactly suited to compressed sensing techniques: we are given an underdetermined linear mapping  $\hat{S}: \mathcal{C} \rightarrow \mathcal{T}$ , and we wish to find

$$\delta \mathbf{c}_0 = \underset{\delta \mathbf{c}}{\operatorname{argmin}} \|\delta \mathbf{c}\|_0, \text{ subject to } \hat{S}(\mathbf{c}_{\text{op}} + \delta \mathbf{c}) = \mathbf{t}_{\text{op}} + \delta \mathbf{t}, \quad (1)$$

where  $\mathbf{c}_{\text{op}}$  ( $\mathbf{t}_{\text{op}}$ ) is the control (target) vector at the working point,  $\delta \mathbf{c}$  ( $\delta \mathbf{t}$ ) is the control (target) change from the working point, and  $\|\cdot\|_0$  is the  $L_0$  pseudonorm<sup>12</sup>, which counts the number of non-zero elements in the vector. This type of constrained optimization is known as *regularization*. Unfortunately, regularization with respect to the  $L_0$  pseudonorm is known to be an NP-hard problem<sup>35</sup>, so finding a solution is not computationally feasible for large systems. Compressed sensing avoids this difficulty by regularizing the problem using the  $L_1$  norm (the sum of the absolute values of the entries of the vector), which is computationally efficient. Regularizing with the  $L_1$  norm tends to achieve sparse solutions, while regularizing with the  $L_2$  norm (the square root of the sum of the squares of the entries of the vector) does

not<sup>11,12</sup>.

Results from compressed sensing show that the control vectors produced by this method are indeed sparse<sup>11,12</sup>; however, the question of locality has not been addressed. We first give an intuitive argument for why CODA should yield solutions that are local, and then provide a numerical demonstration of the fact. Consider a classical scenario in which the dot couplings are purely capacitive. Since the capacitive coupling between a dot and an electrode decays inversely with separation, electrodes that are far away tend to require large voltage changes to achieve the same response as closer electrodes. Minimizing the  $L_1$  norm of the voltage vector  $\delta \mathbf{c}$  suppresses such changes in favor of small voltage changes. In other words, the sparse solutions found by CODA are local.

Figure 2 gives a numerical demonstration of the emergence of local solutions when CODA is applied to an accumulation-mode 8-dot device in a Si/SiGe heterostructure with four capacitively coupled double quantum dot qubits. This design represents a natural extension of a recently demonstrated 4-dot device<sup>37</sup>. Figure 2a shows the device design. We model the device using the semiclassical Thomas-Fermi approximation<sup>38</sup> to compute electron densities and potentials, as described in the Methods. The space of operational targets is defined by the eight dot occupations and the four intra-qubit tunnel rates. Our working point is chosen so there is 1 electron per dot, and energy barriers with maximum heights of 1 meV, corresponding roughly to GHz tunnel rates<sup>39</sup> (see Supplemental Information<sup>36</sup> for details about the simulation parameters).

We then calculate small changes in the dot occupations and tunnel rates when electrode voltages are varied near the working point, obtaining the cross-covariance matrix

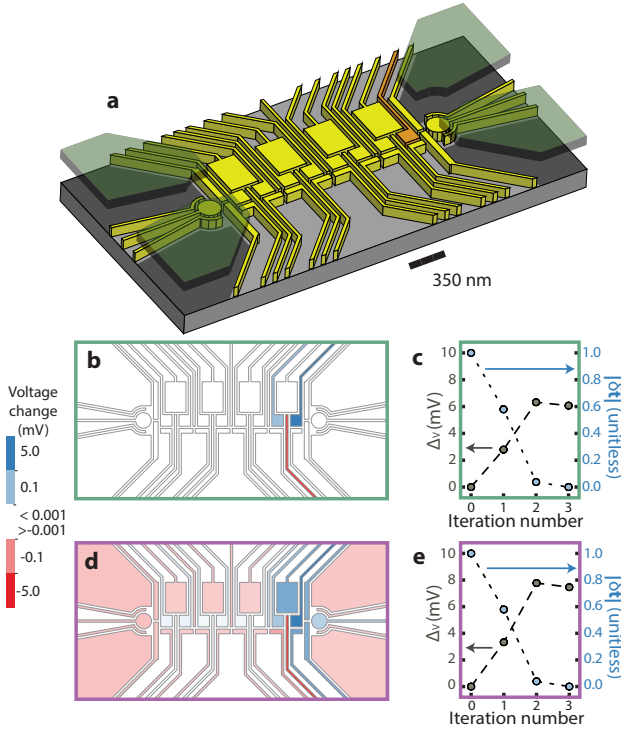


FIG. 2. Locality emerging from the CODA protocol applied to an octuple quantum dot, comprising four double-dot qubits. **a**, Device schematic, with metal gates colored yellow/orange (lower level) and green (upper level). The operational target is to increase the occupation of the right-most quantum dot (underneath the orange electrode) by one electron. **b,d**, Visualization of the voltage variations corresponding to adding one electron to the right-most dot while leaving all other dot occupations and tunnel barriers unchanged, plotted on a logarithmic color scale (electrodes with voltage changes less than  $1 \mu\text{V}$  are colored white). Solutions are obtained by applying the iterative CODA procedure, minimizing the **b**,  $L_1$  and **d**,  $L_2$  norms of the voltage changes found for each cycle. **c,e**, The  $L_1$  norm of the voltage changes relative to the original working point ( $\Delta_V$ ) and the associated distance from the operational target ( $|\delta t|$ , see Supplemental Information<sup>36</sup> for details about this metric) across each iteration of the CODA procedure. Solutions are obtained by minimizing the **c**,  $L_1$  and **e**,  $L_2$  norms of the voltage changes found for each cycle. These results demonstrate that the voltage changes obtained using  $L_1$  minimization are local, requiring adjustment of a small number of electrodes, while many electrode voltages are changed when  $L_2$  minimization is used.

that describes the effective linear response of the system. Once the cross-covariance matrix is determined, we may specify a desired output. Here, we focus on the operational target of adding one electron to the right-most quantum dot, while holding all the other tunnel rates and dot occupations constant.

As this is a linear approximation of a non-linear process, a single cycle of the CODA procedure will not necessarily generate voltages which will yield the desired operational target in the device. Regularization therefore




FIG. 3. Improving double quantum dot designs through CODA. **a**, A schematic of a series of double-dot devices with different SiGe spacer thicknesses. The heterostructure, which is not shown to scale here, consists of 200 nm of  $\text{Si}_{0.7}\text{Ge}_{0.3}$  (the SiGe “buffer”), 10 nm of Si (the Si quantum well), and a layer of  $\text{Si}_{0.7}\text{Ge}_{0.3}$  with variable height  $h$  (the SiGe “spacer”). **b**, The gate geometry used for the devices. The design consists of an upper layer and lower layer of electrodes, depicted in blue and orange, respectively, separated by 80 nm of oxide. The lower layer of electrodes is separated from the SiGe spacer by 10 nm of oxide. **c**, The CODA procedure is applied to a range of devices with different spacer heights,  $h$ . For each simulation we begin with an initial working point corresponding to one electron per dot and an interdot tunnel barrier of 1 meV. We then use CODA to add one electron to the right dot, while keeping the occupation of the left dot and the barrier height fixed. The  $L_1$  norm,  $\Delta_V$ , of the voltage required to achieve this retuning monotonically increases with the height of the SiGe spacer, indicating that thinner spacer layers allow the dots to be controlled more easily.

requires a closed-loop iteration, as indicated in Fig. 1. To reliably converge, the step size needs to be sufficiently small that the linearized model remains approximately correct. In Fig. 2b, we show the voltage changes needed to achieve convergence of the CODA procedure—in this case, after three iterative steps. Note that although voltages are allowed to vary on all the electrodes, only three voltages are found to vary in the  $L_1$ -regularized solution, corresponding to electrodes that are most local to the target of interest. Hence, we confirm that the  $L_1$  norm is an effective proxy for local control, and that CODA is a practical tool for tuning a device, because it picks out control solutions that are both sparse and local while achieving a specified operational target. For comparison, we also performed an alternative optimization procedure based on  $L_2$  regularization, with results shown in Fig. 2d. Note that although the latter procedure achieves the same target objectives, the solution involves voltage changes on almost all of the electrodes, indicating that this solution is neither local nor sparse.

We now show how CODA can be used to optimize device designs, via the objective comparison of different device designs. We focus on a single double-dot qubit design, shown in Fig. 3a, which can be replicated to form

a multi-qubit device. Since CODA automatically localizes the control voltages, using it to design a high-quality unit cell leads naturally to multi-qubit devices that are locally controllable.

We again use the semiclassical Thomas-Fermi approximation<sup>38</sup> to model the devices, with working points similar to those used for the 8-dot device. The heterostructure of the simulated devices is shown in Fig. 3a (see Methods for details). To demonstrate how CODA can be used to compare different devices, we consider several devices with the same gate geometry, but with different SiGe spacer heights. In the device with a 30 nm SiGe spacer, we choose an initial working point that corresponds to 1 electron in each quantum dot, and an interdot tunnel barrier of height 1 meV. We next apply the same CODA procedure used in Fig. 2 to add one electron to the right dot, leaving the occupation of the left dot and the tunnel barrier height unchanged. We then increase the spacer height and repeat the procedure. We note here that each new initial working point, corresponding to a new spacer height, was obtained from the previous initial point by allowing CODA to retune the appropriate electrode voltages, yielding a fully automated procedure.

The results of our CODA investigation, shown in Fig. 3, help to compare and contrast the effectiveness of different device designs. Since  $L_1$  regularization yields solutions that are local and sparse, we expect that smaller values of  $L_1$  will yield devices that are easier to control. In Fig. 3c, we see that reducing the height of the SiGe spacer layer improves the device controllability, as measured by smaller magnitude changes in the control vectors (small  $\Delta_V$ ). This is intuitively reasonable because narrow spacers yield strong capacitive couplings between the dots and the top electrodes. In real devices, the advantages of narrow spacers may be offset by the proximity to disorder in the oxide, which is not included in our simulations. However, it is clear that CODA can be an effective tool for assessing designs, and that many types of design questions can be addressed in this way.

In summary, we have introduced a protocol for the Compressed Optimization of Device Architectures, which applies an  $L_1$  regularization scheme to a linearized model of device operation. We have demonstrated the effectiveness of this scheme by considering its application to semiconductor nanoelectronic quantum dot systems. As devices continue to grow in complexity, such automated control schemes will be essential for design and operation. Our protocol is computationally efficient to implement, and it provides a systematic approach for achieving local and sparse control. Through realistic semiclassical simulations of double-dot devices, we have illustrated how the CODA scheme can be used for quantitative benchmarking and device development. This method provides a path toward the rational design and operation of extensible quantum nanodevices.

## Methods

We perform semi-classical Thomas-Fermi

calculations<sup>38</sup> using the COMSOL Multiphysics software package to solve a nonlinear Poisson equation in three dimensions. We use zero-field boundary conditions on all sides of the simulated domain, with the exception of the bottom of the SiGe buffer, which is grounded. We assume the following heterostructure profile for all the modeled devices, unless otherwise specified. This profile is consistent with the accumulation-mode devices described in refs. 8 and 18: 200 nm of  $\text{Si}_{0.7}\text{Ge}_{0.3}$  (with dielectric constant  $\epsilon = 13.19$ ), a 10 nm Si quantum well ( $\epsilon = 11.7$ ), 30 nm of  $\text{Si}_{0.7}\text{Ge}_{0.3}$ , 10 nm of  $\text{Al}_2\text{O}_3$  ( $\epsilon = 9.0$ ), a 10 nm layer of metallic electrodes, 80 nm of  $\text{Al}_2\text{O}_3$ , and a second 10 nm layer of metallic electrodes.

The occupations of the dots and the leads are found by integrating the induced 2D charge density over the accumulation regions. The heights of the tunnel barriers are calculated by numerically finding the saddle point over the appropriate potential energy landscape calculated by the self-consistent solver. By changing the voltage on the electrodes, the relevant occupations and tunnel barrier heights can be altered to achieve a desired working point, which in this case consists of one electron per dot and tunnel barrier heights of 1 meV. Small ( $\sim 1$  mV) voltage perturbations are then applied to each electrode, and the resulting changes in the operational targets are computed. These data are used to find the mapping in the equation  $\hat{S} : \mathcal{C} \rightarrow \mathcal{T}$  of Eq. 1 (*i.e.*, the cross-covariance matrix).

The  $L_1$  and  $L_2$  regularized optimization procedures are implemented using the CVXPY computational package<sup>40</sup>. For additional simulation details, see the Supplemental Information<sup>36</sup>.

*Acknowledgements*– The authors thank J. B. Aidun, J. Moussa, N. T. Jacobson, R. P. Muller, and C. Tahan for useful comments and discussions. This work was supported in part by the Laboratory Directed Research and Development program at Sandia National Laboratories and ARO (W911NF-12-1-0607 and W911NF-12-R-0012), NSF (PHY-1104660), ONR (N00014-15-1-0029). Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000.

*Author Contributions*– JKG, AF, MF, SNC, RBK, and MAE developed the compressed control scheme. MAE, DRW, and JKG designed the devices. AF and JKG developed and carried out the device simulations. AF, JKG, MF, and SNC wrote the manuscript, with input from all authors.

*Competing financial interests*– JKG, DRW, MAE, MF, and SNC are co-inventors on a patent application related to some of the nanostructure designs described in this work.

*Additional Information*– Supplementary information accompanies this paper. Correspondence and requests for materials should be addressed to John King Gamble

(jkgambl@sandia.gov).

## I. SUPPLEMENTAL INFORMATION

### A. CODA procedure

Here, we provide further details about the CODA protocol sketched in Fig. 1 of the main text. Suppose that we have  $m$  operational targets and  $n$  controls, and that  $n > m$ , so the system is underconstrained. We first identify a working point of experimental interest  $(\mathbf{c}_{\text{op}}, \mathbf{t}_{\text{op}})$  and assume that perturbations of the device from the working point are small. We then linearize the mapping  $\hat{S} : \mathcal{C} \rightarrow \mathcal{T}$  around  $(\mathbf{c}_{\text{op}}, \mathbf{t}_{\text{op}})$ ; we express the resulting linear map as a matrix  $\hat{S}_{\text{op}}$ .

In general, we define

$$\delta \mathbf{c}_\alpha = \underset{\delta \mathbf{c}}{\text{argmin}} \|\delta \mathbf{c}\|_\alpha, \text{ subject to } \hat{S}_{\text{op}}(\mathbf{c}_{\text{op}} + \delta \mathbf{c}) = \mathbf{t}_{\text{op}} + \delta \mathbf{t}, \quad (2)$$

where  $\delta \mathbf{c}$  ( $\delta \mathbf{t}$ ) is the control (target) change from the working point,  $\|\cdot\|_\alpha$  for  $\alpha \geq 1$  is the  $L_n$  norm, and  $\|\cdot\|_0$  is the  $L_0$  pseudonorm<sup>12</sup>, which counts the number of non-zero elements in the vector. Although  $\delta \mathbf{c}_0$  would be sparse by definition, finding  $\delta \mathbf{c}_0$  is computationally infeasible for large systems, so we instead find  $\delta \mathbf{c}_1$ , which tends to be similarly sparse, as shown in the field of compressed sensing<sup>11,12</sup>.

The full protocol for CODA is specified as follows:

1. Identify a working point  $(\mathbf{c}_{\text{op}}, \mathbf{t}_{\text{op}})$  satisfying  $\hat{S}(\mathbf{c}_{\text{op}}) = \mathbf{t}_{\text{op}}$ . In practice,  $\hat{S}$  represents the action of some physical device, and  $\mathbf{t}_{\text{op}}$  is the computed target vector corresponding to the input control  $\mathbf{c}_{\text{op}}$  (e.g. voltages).
2. Consider a set of linearly independent, small control variations  $\{\epsilon_1, \epsilon_2, \dots, \epsilon_n\}$  ( $\epsilon_i \in \mathcal{C}$ ) about the working point. (We took each  $\epsilon_j$  to correspond to a small voltage change on a single electrode.) Perform  $n$  simulations  $\hat{S}(\mathbf{c}_{\text{op}} + \epsilon_i)$  to obtain the resulting  $n$  data points  $\mathbf{t}_{\text{op}} + \delta \mathbf{t}_i$ . Since the control variations were small, by linearity  $\hat{S}_{\text{op}}(\epsilon_i) = \delta \mathbf{t}_i$ .
3. From the collection of  $\{\epsilon_1, \epsilon_2, \dots, \epsilon_n\}$  and the associated  $\{\delta \mathbf{t}_1, \delta \mathbf{t}_2, \dots, \delta \mathbf{t}_n\}$ , construct the linear operator  $\hat{S}_{\text{op}}$  using least squares.
4. For a desired target variation  $\delta \mathbf{t}$ , identify the control variation  $\delta \mathbf{c}_1$  as defined in Eq. (2). Using a convex program (such as the matrix-free cone solver implemented in the CVXPY package<sup>40</sup>), find  $\delta \mathbf{c}_1$  which has minimal  $\|\delta \mathbf{c}\|_1$  subject to the constraint  $\hat{S}_{\text{op}}(\delta \mathbf{c}) = \delta \mathbf{t}$ .
5. Due to the nonlinearity of the system, in general  $\hat{S}(\mathbf{c}_{\text{op}} + \delta \mathbf{c}_1) \neq \mathbf{t}_{\text{op}} + \delta \mathbf{t}$ . However, because the linear approximation is valid within a neighborhood

around the working point, there must exist a scalar  $0 < \gamma \leq 1$  such that  $\hat{S}(\mathbf{c}_{\text{op}} + \gamma \delta \mathbf{c}_1) \approx \mathbf{t}_{\text{op}} + \gamma \delta \mathbf{t}$ . Use a series of simulations to find the maximum value of  $\gamma$  for which this equation holds,  $\gamma_0$ . Define a new control working point  $\mathbf{c}'_{\text{op}} = \mathbf{c}_{\text{op}} + \gamma_0 \delta \mathbf{c}_1$ , and new desired change in target  $\delta \mathbf{t}' = (1 - \gamma_0) \delta \mathbf{t}$ . Repeat this procedure until convergence is achieved at the final target.

### B. Simulation details

We perform semi-classical Thomas-Fermi calculations<sup>38</sup> using the COMSOL Multiphysics software package to solve a nonlinear Poisson equation in three dimensions. Within the Si quantum well, we define a plane of charge with the charge density given by

$$\sigma_{2D}(x, y) = -2 \times 2 \times \frac{em_{\text{eff}}(U(x, y) + E_F)}{2\pi\hbar^2} \times \theta(U(x, y) + E_F), \quad (3)$$

where  $e$  is the charge of an electron,  $m_{\text{eff}} = 0.19 m_{\text{electron}}$  is the transverse effective mass of a conduction electron in silicon,  $U(x, y)$  is the strength of the electrostatic potential energy as a function of position,  $E_F$  is the Fermi energy (which we take to be at ground), and  $\theta(x)$  is the step function. The two prefactors account for the spin and valley degeneracies.

The details of the working point used in the analysis of the 8-dot device are given in the supplemental file `8DotDevice.txt`. In this file, the physical attributes are listed first. The dot occupations are expressed in electron numbers, and the tunnel barrier heights in eV. Voltages are given for each electrode, with the following labeling convention defined with respect to Fig. 2a of the main text. Beginning with the upper layer of electrodes, Electrode 1 is in the lower-right corner of the schematic, and the ordering proceeds clockwise. In the lower layer of gates, Electrode 5 is in the lower-right corner, and the ordering again proceeds clockwise.

Similar details for working points on the devices shown in Fig. 3 are given in the supplemental file `VaryingHeight.txt`. As before, the dot occupations are given in numbers of electrons and the tunnel barrier heights in eV. Here, the labeling convention for the electrodes begins with the upper layer at the electrode in the upper right corner of the schematic and proceeds clockwise. On the lower layer of electrodes, the labeling begins at the electrode in the upper-left corner and proceeds clockwise.

Our CODA procedure requires all the components of the control vector to have the same units (and comparable magnitudes for numerical stability). The operational targets considered in our simulations were electron occupations and tunnel barrier heights. To make these quantities comparable for the devices studied here, we multiplied the tunnel barrier height by a factor of 10, to ensure rapid convergence.

- <sup>1</sup> Petta, J. R. *et al.* Coherent manipulation of coupled electron spins in semiconductor quantum dots. *Science* **309**, 2180 (2005).
- <sup>2</sup> Koppens, F. *et al.* Control and detection of singlet-triplet mixing in a random nuclear field. *Science* **309**, 1346–1350 (2005).
- <sup>3</sup> Prance, J. R. *et al.* Single-shot measurement of triplet-singlet relaxation in a Si/SiGe double quantum dot. *Phys. Rev. Lett.* **108**, 046808 (2012).
- <sup>4</sup> Maune, B. M. *et al.* Coherent singlet-triplet oscillations in a silicon-based double quantum dot. *Nature* **481**, 344–347 (2012).
- <sup>5</sup> Shulman, M. D. *et al.* Demonstration of entanglement of electrostatically coupled singlet-triplet qubits. *Science* **336**, 202 (2012).
- <sup>6</sup> Pla, J. J. *et al.* A single-atom electron spin qubit in silicon. *Nature* **489**, 541–545 (2012).
- <sup>7</sup> Kim, D. *et al.* Quantum control and process tomography of a semiconductor quantum dot hybrid qubit. *Nature* **511**, 70–74 (2014).
- <sup>8</sup> Wu, X. *et al.* Two-axis control of a singlet-triplet qubit with an integrated micromagnet. *P. N. A. S.* **111**, 11938–11942 (2014).
- <sup>9</sup> Fuechsle, M. *et al.* A single-atom transistor. *Nature Nano.* **7**, 242–246 (2012).
- <sup>10</sup> Koza, J. R., Bennett, F. H., Andre, D., Keane, M. A. & Dunlap, F. Automated synthesis of analog electrical circuits by means of genetic programming. *IEEE Transactions on Evolutionary Computation* **1**, 109–128 (1997).
- <sup>11</sup> Candès, E. J., Romberg, J. K. & Tao, T. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics* **59**, 1207–1223 (2006).
- <sup>12</sup> Donoho, D. Compressed sensing. *IEEE Transactions on Information Theory* **52**, 1289–1306 (2006).
- <sup>13</sup> Hanson, R., Kouwenhoven, L. P., Petta, J. R., Tarucha, S. & Vandersypen, L. M. K. Spins in few-electron quantum dots. *Rev. Mod. Phys.* **79**, 1217–1265 (2007).
- <sup>14</sup> Loss, D. & DiVincenzo, D. P. Quantum computation with quantum dots. *Phys. Rev. A* **57**, 120–126 (1998).
- <sup>15</sup> Kane, B. E. A silicon-based nuclear spin quantum computer. *Nature* **393**, 133–137 (1998).
- <sup>16</sup> Medford, J. *et al.* Self-consistent measurement and state tomography of an exchange-only spin qubit. *Nature Nano.* **8**, 654–659 (2013).
- <sup>17</sup> Medford, J. *et al.* Quantum-dot-based resonant exchange qubit. *Phys. Rev. Lett.* **111**, 050501 (2013).
- <sup>18</sup> Kawakami, E. *et al.* Electrical control of a long-lived spin qubit in a si/sige quantum dot. *Nature Nano.* **9**, 666–670 (2014).
- <sup>19</sup> Veldhorst, M. *et al.* An addressable quantum dot qubit with fault-tolerant control fidelity. *Nature Nano.* **9**, 981–985 (2014).
- <sup>20</sup> Eng, K. *et al.* Isotopically enhanced triple-quantum-dot qubit. *Science Advances* **1**, e1500214 (2015).
- <sup>21</sup> Kim, D. *et al.* Microwave-driven coherent operation of a semiconductor quantum dot charge qubit. *Nature Nano.* **10**, 243–247 (2015).
- <sup>22</sup> Kim, D. *et al.* High-fidelity resonant gating of a silicon-based quantum dot hybrid qubit. *Npj Quantum Information* **1**, 15004 (2015).
- <sup>23</sup> Delbecq, M. R. *et al.* Quantum dephasing in a gated gaas triple quantum dot due to nonergodic noise. *Phys. Rev. Lett.* **116**, 046802 (2016).
- <sup>24</sup> Takeda, K. *et al.* A fault-tolerant addressable spin qubit in a natural silicon quantum dot (2016). Preprint at <http://arxiv.org/abs/1602.07833>.
- <sup>25</sup> Martins, F. *et al.* Noise suppression using symmetric exchange gates in spin qubits. *Phys. Rev. Lett.* **116**, 116801 (2016).
- <sup>26</sup> Pla, J. J. *et al.* High-fidelity readout and control of a nuclear spin qubit in silicon. *Nature* **496**, 334–338 (2013).
- <sup>27</sup> Muhonen, J. T. *et al.* Storing quantum information for 30 seconds in a nanoelectronic device. *Nature Nano.* **9**, 986–991 (2014).
- <sup>28</sup> Zwanenburg, F. A. *et al.* Silicon quantum electronics. *Rev. Mod. Phys.* **85**, 961 (2013).
- <sup>29</sup> Harvey-Collard, P. *et al.* Nuclear-driven electron spin rotations in a single donor coupled to a silicon quantum dot (2015). Preprint at <http://arxiv.org/abs/1512.01606>.
- <sup>30</sup> Muhonen, J. T. *et al.* Quantifying the quantum gate fidelity of single-atom spin qubits in silicon by randomized benchmarking. *Journal of Physics: Condensed Matter* **27**, 154205 (2015).
- <sup>31</sup> Laucht, A. *et al.* Electrically controlling single-spin qubits in a continuous microwave field. *Science Advances* **1**, e1500022 (2015).
- <sup>32</sup> Kelly, J. *et al.* Optimal quantum control using randomized benchmarking. *Phys. Rev. Lett.* **112**, 240504 (2014).
- <sup>33</sup> Shulman, M. D. *et al.* Suppressing qubit dephasing using real-time Hamiltonian estimation. *Nat. Commun.* **5**, 5156 (2014).
- <sup>34</sup> Baart, T. A., Eendebak, P. T., Reichl, C., Wegscheider, W. & Vandersypen, L. M. K. Computer-automated tuning of semiconductor double quantum dots into the single-electron regime. *Applied Physics Letters* **108**, 213104 (2016).
- <sup>35</sup> Natarajan, B. K. Sparse approximate solutions to linear systems. *SIAM Journal on Computing* **24**, 227–234 (1995).
- <sup>36</sup> Frees, A. *et al.* Supplemental information. *Supplemental Information* (2016).
- <sup>37</sup> Ward, D. *et al.* State-conditional coherent charge qubit oscillations in a si/sige quadruple quantum dot. *arXiv preprint arXiv:1604.07956* (2016).
- <sup>38</sup> Stopa, M. Quantum dot self-consistent electronic structure and the coulomb blockade. *Phys. Rev. B* **54**, 13767–13783 (1996).
- <sup>39</sup> Shirkhorshidian, A. *et al.* Transport spectroscopy of low disorder silicon tunnel barriers with and without Sb implants. *Nanotechnology* **26**, 205703 (2015).
- <sup>40</sup> Diamond, S. & Boyd, S. Cvxpy: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* **17**, 1–5 (2016).